Remarks

Claims 1-16 remain in the application.

Claim Rejections -- 35 USC 102

As a preliminary matter, applicants respectfully submit that claims 10, 12, 14 and 16 were previously amended as of October 19, 2007 to expressly recite that the link-loss-learn protocol involves "upon detecting a link failure at the port, the MAC address table is cleared of all MAC address entries therein" or language to that effect. Hence, applicants respectfully submit that this suggestion of the Examiner (under section 3 of the latest office action) has already been implemented.

Claims 1, 10 and 14 stand rejected under 35 USC 102(e) as being anticipated by Ishii (US 2001/0021177 A1). Applicants respectfully traverse this rejection.

Let us consider claim 1, which is reproduced below for convenience of reference.

1. A method of fault recovery by a switch in a local area network, the method comprising:

   detecting a link failure at a port <u>of the switch</u>; and

   clearing all medium access control (MAC) address entries from a MAC address table <u>of the switch in response to the link failure detection</u> and **<u>without</u> receiving from outside the switch any signal that signifies that the MAC address table of the switch is to be cleared.**

(Emphasis added.)

Now consider the contention in the latest rejection that Fig. 1, table 9 and paragraphs 184-185 of Ishii disclose the limitation of claim 1 which states,

"**clearing all medium access control (MAC) address entries from a MAC address table <u>of the switch in response to the link failure detection</u>** and without receiving from outside the switch any signal that signifies that the MAC address table of the switch is to be cleared." (Emphasis added.) Applicants respectfully submit that "the switch" in this limitation clearly refers to the **same** switch as in the previous limitation by virtue of its *antecedent* basis. In other words, claim 1 clearly requires that both detecting the link failure and clearing the MAC address table are performed by and within a **single** switch.

Below is the text of paragraphs 184-185 which were cited in the latest office action.

> [0184] When a forwarding table clearing portion 10 detects a TC detection flag set by the topology change detecting portion 5 , it deletes the contents (database information) of the forwarding table 9 .
>
> [0185] In more detail, when a BPDU in which a TC detection flag is set **is received from a root bridge** by means of the topology change detecting portion 5 , database information contained in the forwarding table 9 is deleted.

(Emphasis added.)

As seen above, Ishii clearly teaches deleting contents the forwarding table 9 "when a BPDU in which a TC detection flag is set **is received from a root bridge** by means of the topology change detecting portion 5 ...." (Emphasis added.) In other words, Ishii teaches deleting contents of a forwarding table in a particular bridge upon receiving a particular packet (BPDU in which a TC detection flag is set) from a **separate** bridge (the root bridge). As such, this disclosure clearly <u>cannot</u> read on claim 1 because it involves <u>two</u> separate bridges (the receiving bridge and the root bridge) while claim 1 clearly is performed by and within a <u>single</u> switch.

Moreover, claim 1 was previously amended so as to explicitly recite that the MAC address table is cleared "... **<u>without</u> receiving from outside the switch any signal that signifies that the MAC address table of the switch is to be cleared**." In contrast, Ishii clearly discloses that a special packet, namely

a BPDU packet with a TC (topology change) flag set, **is received**, and that this special packet is taught as signifying to all receiving bridges that their forwarding tables are to be cleared. Hence, it is abundantly clear in the express language of claim 1 that Ishii cannot read on this claim.

Applicants further note that the present application discusses the STP method using topology change notifications in the form of a BPDU (the same or similar to Ishii) on page 4, line 4 through page 5, line 19 in relation to FIG. 2. As stated therein, "The STP method is considered to be clever because traffic related to entries not affected by the broken link continues to be transmitted and those unaffected entries in the MAC address table do not have to be relearned." (Page 5, lines 12-14.) However, as further discussed in the specification, applicants have come up with the claimed ***counter-intuitive*** solution of automatically clearing the entire table and forcing all addresses to be relearned, whether or not the addresses are affected by the link failure. A ***surprising*** result of this counter-intuitive solution is that it provides for more rapid fault recovery, ***despite*** the fact that unaffected entries in the MAC address table must be relearned.

This innovative feature, dubbed Link-Loss Learn® by applicants, has been implemented as a successful, distinctive product feature on certain of its Garrettcom® Magnum® switches. For example, attached is a marketing document entitled, "Link-Loss-Learn on Magnum Managed Ethernet Switches … a Technical Brief" which discusses this product feature. Applicants respectfully submit that the ***commercial success*** of these Garrettcom® Magnum® switches implementing the Link-Loss Learn® feature provides evidence of ***secondary considerations*** that **rebut** rejections of unpatentability under 35 USC 103.

For at least the above-discussed reasons, applicants respectfully submit that claim 1 overcomes its rejection.

Similar to claim 1, claim 10 clearly recites that the MAC address table of a network apparatus is cleared when a link failure is detected at a port of the **same** apparatus. In contrast, as discussed above, Ishii teaches deleting contents of a

forwarding table in a particular bridge upon receiving a particular packet (BPDU in which a TC detection flag is set) from a **separate** bridge (the root bridge). In addition, claim 10 also expressly states that the address table is cleared "**without receiving from outside the apparatus any signal that signifies that the MAC address table of the apparatus is to be cleared**." Hence, claim 10 also overcomes this rejection.

Also similar to claim 1, claim 14 clearly recites a switch implementing a link-loss-learn protocol which involves "upon detecting a link failure at a port **of the switch**, clearing a media access control (MAC) address table of all MAC entries therein, **without receiving from outside the switch any signal that signifies that the MAC address table of the switch is to be cleared**." Hence, claim 14 also overcomes this rejection.

Thus, for the reasons discussed above, applicants respectfully submit that claims 1, 10 and 14 each overcome this rejection.

Claim Rejections -- 35 USC 103

Claims 2-4, 7, 11-13, 15 and 16 stand rejected under 35 USC 103 as being unpatentable over Bare (US 2003/0016624) in view of Ishii. This rejection is respectfully traversed.

As discussed above, independent claims 1, 10 and 14 are patentably distinguished over Ishii. The citation to Bare does not cure the deficiencies discussed above in regard to Ishii. Furthermore, applicants respectfully submit that the ***commercial success*** of the Garrettcom® Magnum® switches implementing the Link-Loss Learn® feature provides evidence of ***secondary considerations*** that **rebut** rejections of unpatentability under 35 USC 103. Therefore, applicants respectfully submit that dependent claims 2-4, 7, 11-13, 15 and 16 overcome this rejection.

Claims 8--9 also stand rejected under 35 USC 103 as being unpatentable over Bare (US 2003/0016624) in view of Ishii. This rejection is respectfully traversed.

As discussed above, independent claim 1 is patentably distinguished over Ishii. The citation to Bare does not cure the deficiencies discussed above in regard to Ishii. Furthermore, applicants respectfully submit that the ***commercial success*** of the Garrettcom® Magnum® switches implementing the Link-Loss Learn® feature provides evidence of ***secondary considerations*** that **rebut** rejections of unpatentability under 35 USC 103. Therefore, applicants respectfully submit that dependent claims 8-9 overcome this rejection.

Claim 5 stands rejected under 35 USC 103 as being unpatentable over Bare (US 2003/0016624) in view of Ishii and further in view of Eisen et al. This rejection is respectfully traversed.

As discussed above, independent claim 1 is patentably distinguished over Ishii. The citations to Bare and Eisen et al. do not cure the deficiencies discussed above in regard to Ishii. Furthermore, applicants respectfully submit that the ***commercial success*** of the Garrettcom® Magnum® switches implementing the Link-Loss Learn® feature provides evidence of ***secondary considerations*** that **rebut** rejections of unpatentability under 35 USC 103. Therefore, applicants respectfully submit that dependent claim 5 overcomes this rejection.

Lastly, claim 6 stands rejected under 35 USC 103 as being unpatentable over Bare (US 2003/0016624) in view of Ishii and further in view of Tanoue. This rejection is respectfully traversed.

As discussed above, independent claim 1 is patentably distinguished over Ishii. The citations to Bare and Tanoue do not cure the deficiencies discussed above in regard to Ishii. Furthermore, applicants respectfully submit that the ***commercial success*** of the Garrettcom® Magnum® switches implementing the Link-Loss Learn® feature provides evidence of ***secondary considerations*** that **rebut** rejections of unpatentability under 35 USC 103. Therefore, applicants respectfully submit that dependent claim 6 overcomes this rejection.

Conclusion

For the above-discussed reasons, applicants respectfully submit that the claims are clearly shown to be patentably distinguished over the cited art. Favorable action is respectfully solicited.

Allowance of the pending claims is respectfully requested as this is the fifth office action to which applicants have responded for this application. Applicants respectfully submit that the pending claims have already received a very thorough examination as evidenced by the numerous references which have been cited and patentably distinguished by the applicants in the course of responding to the five office actions.

Finally, applicants note that a petition to make special was granted for this application based on the age of an applicant. Yet, this application has now been pending at the patent office for about 5 ½ years. For this additional reason, applicants respectfully request a speedy resolution to this examination.

If for any reason an insufficient fee has been paid, the Commissioner is hereby authorized to charge the insufficiency to Deposit Account No. 50-2427 of Okamoto & Benedicto LLP.

Respectfully Submitted,

Dated: April 1, 2009

/James K. Okamoto, Reg. No. 40,110/
James K. Okamoto, Reg. No. 40,110
Okamoto & Benedicto LLP
P.O.Box 641330
San Jose, CA 95164-1330
Tel: (408) 436-2111
Fax: (408) 436-2114

Link-Loss-Learn on Magnum Managed Ethernet Switches . . . a Technical Brief

GarrettCom
Ethernet at its Best

Introduction

GarrettCom's Link-Loss-Learn™ feature, designed to simplify and speed up recovery for Ethernet switches used in redundant LAN topologies, is implemented in the Magnum™ Managed Ethernet Switches and ES42 unmanaged switches. It addresses the time delay associated with changing the network addresses that are normally stored in a switch's memory. With Link-Loss-Learn (patent pending), Magnum 6Ks, mP62s, and ES42s are better able to handle some fault recovery situations, and they may improve network reliability and provide faster fault recovery accordingly.

The managed Magnum mP62 Hardened Ethernet Switch, designed for "edge" applications where data devices connect into the LAN, was the first Magnum product to offer the Link-Loss-Learn feature. The unmanaged ES42 designed for "edge" applications where data devices connect into the LAN, has Port 1 and Port 2 set for LLL operation as a factory default setting.

The Link-Loss-Learn feature is placed into operation in a Magnum Managed Switch under the users control. It is not "automatic," but rather it is turned on with software commands by the user at initial set-up of the Magnum MNS management software. If Link-Loss-Learn is not turned on, the Magnum Switch operates and functions the same as an ordinary Ethernet Switch, learning the MAC addresses of attached nodes and retaining those addresses in memory until they eventually age out or power is turned off. The Link-Loss-Learn feature will improve fault recovery in ring topologies, but it can never hurt by going into operation unexpectedly.

Users may enable Link-Loss-Learn on a port-by-port basis in managed switches. A typical configuration selection enables Link-Loss-Learn on the Ring Member ports (usually two fiber ports) since these ports are normally used to connect into the redundant network upstream. However, the Link-Loss-Learn feature may be enabled on any combination of copper and fiber ports of any speed - - or it can be totally inactive.

How does Link-Loss-Learn assist with Fault Recovery?

When a LAN is functioning normally, the LINK indicator is present for each port in use. A fault (a cable cut, or a unit losing power, or a unit failing while in operation) will usually cause LINK to be lost on one or more ports. Therefore, the loss of LINK during normal operation is interpreted as a signal that something has gone wrong, and in a redundant LAN, recovery operations must be brought into action to restore Ethernet traffic to its expected performance level.

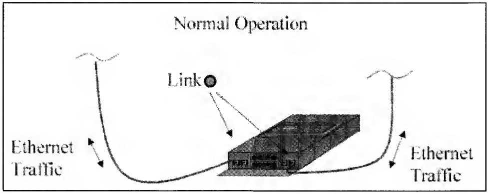

When LINK fails on a port in a redundant LAN, another back-up port is expected to take over and keep the network packets flowing. The back-up port is connected and ready to provide service. However, a normal Ethernet Switch engine will continue to use its old address table, and will continue to try to

forward packets to the failed port. This will go on until the address table aging time expires for the addresses whose connection was lost, which can be as much as 4 or 5 minutes.

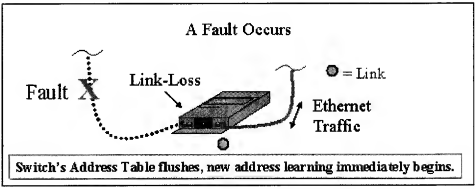

When standard 802.1d Spanning Tree Protocol is implemented, once a topology change is detected, the STP aging timer is set to 15 seconds until the topology is re-computed and / or reconfigured. The process of re-computing as well as reconfiguring the LAN can take equally as long, even in a simple set-up. (Note - complex set-ups such as multi-level meshes take much longer). For some industrial networks, this time of less than a minute for fault recovery is an acceptable delay, and standard Spanning Tree is an acceptable solution. For faster fault recovery and restoration times, Link-Loss-Learn can help.

The Link-Loss-Learn feature improves the recovery time by forcing the Magnum Switch's address table to be flushed when LINK is lost on any designated port. The effect on the operation of the Switch is the same as upon power up. The first packet is broadcasted and its address is learned. This continues rapidly until all addresses are learned and operation is normal . . . but with new information now in the address table on how to switch packets. Some bandwidth is used unnecessarily during the re-learning, but the recovery process is not delayed. Thus, the immediate re-learning of the addresses of attached devices results in fast re-routing of the network traffic passing through the Switch.

Because the Link-Loss-Learn feature is very fast (it takes just a few milliseconds), the Magnum Switch will not be the gating item for fault recovery in a redundant LAN. Whether the redundant paths upstream are controlled by 802.1d Standard Spanning Tree Protocol (STP), or by 802.1s Tagged VLAN Spanning Tree Protocol, or by 802.1w Rapid Spanning Tree Protocol, or manually such as in a bench-test situation, Magnum Switches with Link-Loss-Learn can reset their address table and participate in the LAN configuration change and network recovery faster than the other Ethernet elements.

For redundant systems, simplicity is a virtue. Redundant LANs and fail-over scenarios are necessarily complicated, but complexity can also add to the risk that not all will go well when the critical time arrives.

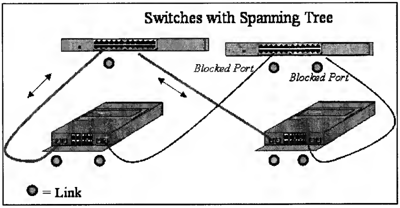

It is better to keep things as simple as possible and minimize complexity. The Link-Loss-Learn feature fits in with this philosophy. Typically, a Magnum Switch (such as an mP62, illustrated here) with LLL is used as an edge switch in a redundant LAN configuration. While an mP62 can run Spanning Tree and can participate in failure recovery schemes accordingly, it can also perform its role in a redundant LAN recovery scenario via the Link-Loss-Learn feature simply

and independently of other things that are going on. One less thing to go wrong. Many applications are well served with the Magnum Edge Switch running the simple Link-Loss-Learn feature rather than a managed switch running the more complex Spanning Tree Protocol or Rapid Spanning Tree Protocol. Magnum products capabilities, with Spanning Tree and Rapid Spanning Tree and the Link-Loss-Learn feature available on many different models and price points, allows the user to choose.

Rings and Strings, and Link-Loss-Learn with Propagation

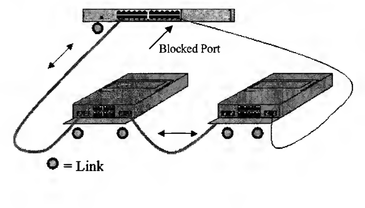

Frequently, a redundant system design using Magnum Switches as edge devices has the Magnums deployed over a distance and interconnected using the fiber ports hooked up in a string, i.e., daisy chained from one switch to the next with fiber media. It is common to continue from the last unit in a string, connecting the 2nd ring member port of the last switch in a given string back into a redundant LAN connection, thus forming a "ring" of switch units. Such a ring must have a port somewhere in the series operating blocked (i.e., not passing packets) so that a correct Ethernet topology exists. The Spanning Tree or similar logic controls which port is blocked in order to manage operating the network and to facilitate recovery from faults.

A Magnum Switch in a string or ring could experience a link loss on a ring member port, indicating a fault has occurred and the string has been broken. Recovery action needs to take place, and a Magnum Switch with the Link-Loss-Learn feature enabled would immediately dump its address table and be ready to relearn addresses and operate in a changed network configuration. But, what about the other switches in the string or ring . . . how will they know to do the same thing?

Switches with Spanning Tree

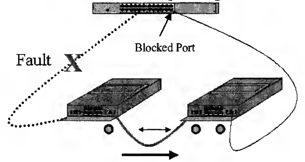

Link-Loss Propagation #1

For this reason, the Link-Loss-Learn feature in Magnum Switches includes a "propagation" function that, upon Link loss on one enabled (ring member) port, temporarily drops Link on any other Link-Loss-Learn enabled ports to propagate the action in the units in the string or ring.

Switches with Spanning Tree

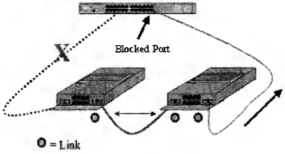

Link-Loss Propagation #2

The Propagation function associated with the Magnum Link-Loss-Learn feature is always present, and its operation only affects those Magnum Switch ports that have the Link-Loss-Learn feature enabled. Users have control over the propagation accordingly, and can control their redundant LAN set-up with Magnum Switches to best suit their application.

Summary

Using the Link-Loss-Learn feature with Propagation, Magnum Managed Ethernet Switches with MNS software, and unmanaged ES42 Edge Switches with LLL, can simplify and speed up recovery from faults in redundant Ethernet LAN configurations. The Link-Loss-Learn feature is applicable to both mesh and string or ring topologies. Typically, using the simple Link-Loss-Learn feature will be better than running Spanning Tree or Rapid Spanning Tree on every Switch in a redundant LAN, increasing reliability by reducing complexity without compromising fault recovery performance.

GarrettCom

Ethernet at its Best

GarrettCom, Inc.
47823 Westinghouse Drive • Fremont, CA 94539 • PH: (510) 438-9071 • FAX: (510) 438-9072
Email: mktg@garrettcom.com • Web: www.GarrettCom.com